arXiv:1412.0845v1 [cs.GT] 2 Dec 2014

On the Robustness of the Approximate Price of Anarchy in Generalized Congestion Games (Full Version)*

Vittorio Bilò †

December 3, 2014

Abstract

One of the main results shown through Roughgarden's notions of smooth games and robust price of anarchy is that, for any sum-bounded utilitarian social function, the worst-case price of anarchy of coarse correlated equilibria coincides with that of pure Nash equilibria in the class of weighted congestion games with non-negative and non-decreasing latency functions and that such a value can always be derived through the, so called, smoothness argument. We significantly extend this result by proving that, for a variety of (even non-sum-bounded) utilitarian and egalitarian social functions and for a broad generalization of the class of weighted congestion games with non-negative (and possibly decreasing) latency functions, the worst-case price of anarchy of ϵ-approximate coarse correlated equilibria still coincides with that of ϵ-approximate pure Nash equilibria, for any $\epsilon \geq 0$. As a byproduct of our proof, it also follows that such a value can always be determined by making use of the primal-dual method we introduced in a previous work. It is important to note that our scenario of investigation is beyond the scope of application of the robust price of anarchy (for as it is currently defined), so that our result seems unlikely to be alternatively proved via the smoothness framework.

1 Introduction

The celebrated notion of robust price of anarchy introduced by Roughgarden in [19, 20] has lately arouse much interest in the determination of inefficiency bounds for pure Nash equilibria which may automatically extend to some of their appealing generalizations, such as mixed Nash equilibria, correlated equilibria and coarse correlated equilibria. These three types of solutions have a particular flavor since, differently from pure Nash equilibria, they are always guaranteed to exist by Nash's Theorem [16][1]; moreover, the last two ones can also be efficiently computed and even easily learned when a game is repeatedly played over time.

*This work was partially supported by the PRIN 2010–2011 research project ARS TechnoMedia: "Algorithmics for Social Technological Networks" funded by the Italian Ministry of University.

†Department of Mathematics and Physics "Ennio De Giorgi", University of Salento, Provinciale Lecce-Arnesano, P.O. Box 193, 73100 Lecce - Italy, Email: vittorio.bilo@unisalento.it.

[1]To this aim, we recall that the set of coarse correlated equilibria contains that of correlated equilibria, which contains that of mixed Nash equilibria, which contains that of pure Nash equilibria.

To this aim, Roughgarden [19, 20] identifies a class of games, called *smooth games*, for which a simple three-line proof, called *smoothness argument*, shows significant upper bounds on the price of anarchy of pure Nash equilibria as long as the social function measuring the quality of any strategy profile in the game is ***sum-bounded***, that is, upper bounded by the sum of the players' costs[2]. He then defines the *robust price of anarchy* of a smooth game as the best-possible (i.e., the lowest) upper bound which can be derived by making use of this argument and provides an *extension theorem* which shows that, still for sum-bounded social functions, the price of anarchy of coarse correlated equilibria of any smooth game is upper bounded by its robust price of anarchy. Finally, he shows that several games considered in the literature happen to be smooth and that the class of (unweighted) congestion games with non-negative and ***non-decreasing latency functions*** is *tight* for the ***utilitarian social function*** (that is, the social function defined as the sum of the players' costs), in the sense that, in this class of games, the worst-case price of anarchy of pure Nash equilibria exactly matches the robust price of anarchy. This last result has been subsequently extended to the class of weighted congestion games by Bhawalkar, Gairing and Roughgarden in [3].

1.1 Our Contribution and Significance

In this work, we generalize the tightness result by Bhawalkar, Gairing and Roughgarden along the following four directions (see Section 2 for formal definitions):

1. the class of games we consider is a broad generalization of that of weighted congestion games. In particular, we focus on *generalized weighted congestion games*, that is, games in which each player's *perceived cost* is defined as a certain linear combination of all the players' *individual costs* originally experienced in some underlying weighted congestion game. Thus, it is quite easy to figure out that the class of generalized weighted congestion games widely extends that of weighted congestion games;

2. the families of social functions we consider are generalizations of both the utilitarian and the egalitarian social functions (where the egalitarian social function is defined as the maximum of the players' costs). In particular, a family of utilitarian social functions is obtained by summing up a certain contribution from each player, whereas a family of egalitarian social functions is obtained by taking the maximum contribution among the players, where each player's contribution is given by a conic combination of the players' individual costs. We stress that such a combination may significantly differ from the one used to define the players' perceived costs, so that there exist social functions in both families that may not be sum-bounded;

3. the latency functions we consider in the definition of the players' individual costs are selected from a family of allowable non-negative functions with no additional restrictions. This permits us to encompass also latency functions not considered so far in the previous tightness results known in the literature, such as, for instance, the widely used fair cost sharing rule induced by the Shapley value [21];

4. the solution concepts we consider are the approximate versions of all the four types of equilibria named so far. In particular, for any real value $\epsilon \geq 0$, we focus on either ϵ-approximate

[2]Throughout the paper, we implicitly assume that all games under consideration are cost minimization ones. All the claimed properties and results can be applied *mutatis mutandis* to the case of payoff maximization games.

pure Nash equilibria and ϵ-approximate coarse correlated equilibria. For the special case of $\epsilon = 0$, one reobtains the notions of pure Nash equilibria and coarse correlated equilibria, so that results for these solution concepts can be obtained as a special case of the ones holding for their approximate versions.

More precisely, but still informally speaking, we prove the following result (Theorem 1 in Section 3):

> *for a variety of utilitarian and egalitarian social functions and for any real value $\epsilon \geq 0$, the worst-case price of anarchy of ϵ-approximate pure Nash equilibria coincides with that of ϵ-approximate coarse correlated equilibria in the class generalized weighted congestion games with non-negative latency functions.*

As it can be appreciated, the above tightness result generalizes the previous one by Bhawalkar, Gairing and Roughgarden along all four directions simultaneously. The technique we use to prove the theorem is the primal-dual method that we introduced in [4]. In fact, as a byproduct of our proof, it also follows that, in the above considered scenario of investigation,

> *the worst-case price of anarchy of ϵ-approximate pure Nash equilibria can always be determined through the primal-dual method.*

We would like to stress that, when adopting the social functions described at point 2, generalized weighted congestion games are not smooth games in general, so that the above tightness result seems unlikely to be reproved via smoothness arguments, at least in the way in which they have been defined and used so far in the literature. This seems to provide an evidence that the primal-dual method may be more powerful than the smoothness framework as far as we focus on congestion games and some of their possible generalizations.

1.2 Related Work

The notion of price of anarchy as a measure of the inefficiency caused by selfish behavior in non-cooperative games has been introduced in a seminal paper by Koutsoupias and Papadimitriou [14] in 1999. Since then, several classes of games have been studied under this perspective. Among these classes, congestion games introduced by Rosenthal in [18] and their weighted variants [15] occupy a preeminent role.

Awerbuch, Azar and Epstein [2] and Christodoulou and Koutsoupias [10] focus on the worst-case price of anarchy of pure Nash equilibria in either weighted and unweighted congestion games under the utilitarian social function. They independently give tight bounds for the case of affine latency functions and almost tight upper and lower bounds for the case of polynomial latency functions with non-negative coefficients. Such a gap has been subsequently closed by Aland et al. in [1]. Moreover, Christodoulou, Koutsoupias and Spirakis [12] obtain tight bounds on the worst-case price of anarchy of approximate pure Nash equilibria in unweighted congestion games for the case of polynomial latency functions with non-negative coefficients, while Christodoulou and Koutsoupias [11] show that the worst-case price of anarchy of correlated equilibria is the same as that for pure Nash equilibria in weighted and unweighted congestion games when considering affine latency functions. As already said, such an equivalence has been further extended to coarse correlated equilibria and to any class of non-negative and non-decreasing latency functions by Roughgarden [19, 20] in the unweighted case and by Bhawalkar, Gairing and Roughgarden [3] in the weighted case, by making use of the smoothness argument and the robust price of anarchy.

Robust bounds on the worst-case price of anarchy have been lately achieved via extensions of the smoothness argument in some generalizations of (unweighted) congestion games. In particular, de Keijzer et al. [13] and Rahn and Schäfer [17] consider the altruistic extension of congestion games in which, similarly to our model of generalized congestion games, the perceived cost of each player is defined as a linear combination of the individual costs of all the players in the game. Anyway, while we do not impose any kind of restriction on such a combination, they consider the case in which the multiplicative coefficients lie in the interval $[0, 1]$ and, for each player i, the contribution of the individual cost of player i to her perceived cost has to be always multiplied by 1. Moreover, they restrict their analysis to the case in which the social function is the sum of the players' individual costs.

Much less attention has been devoted in the literature to the egalitarian social function, for which Christodoulou and Koutsoupias [10] give an asymptotically tight bound on the worst-case price of anarchy in unweighted congestion games with affine latency functions.

We introduced the primal-dual method in [4] as a tool for obtaining tight bounds on the inefficiencies caused by selfish behavior in weighted congestion games and their possible generalizations for a variety of solutions concepts. In particular the primal-dual method has been applied by Bilò, Flammini and Gallotti [7] to derive tight bounds on the worst-case price of anarchy of pure Nash equilibria in congestion games with affine latency functions under the assumption that the players' knowledge is restricted by the presence of an underlying social knowledge graph; by Bilò [5] to derive tight bounds on the worst-case price of stability of pure Nash equilibria in congestion games with affine latency functions and altruistic players; by Bilò and Paladini [9] to derive tight bounds on the approximation ratio of the solutions achieved after a one-round walk of ϵ-approximate best-responses starting from any initial strategy profile in cut games, for any $\epsilon \geq 0$; by Bilò et al. [8] to derive a surprising matching lower bound on the price of anarchy of subgame perfect equilibria in sequential cut games; and by Bilò, Fanelli and Moscardelli [6] to derive significant upper bounds on the price of anarchy of lookahead equilibria in congestion games with affine latency functions.

1.3 Paper Organization

The paper is organized as follows. In the next section, we give all necessary definitions and notation and provide also some preliminary remarks. Section 3 contains the technical contribution of the paper, with the proof of our main theorem. In the last section, we conclude and discuss open problems.

2 Definitions, Notation and Preliminaries

A *weighted congestion game* is a tuple $\mathsf{CG} = \left([n], (w_i)_{i\in[n]}, E, (\Sigma_i)_{i\in[n]}, (\ell_e)_{e\in E}\right)$ such that $[n] = \{1, 2, \ldots, n\}$ is a set of $n \geq 2$ players, $w_i > 0$ is the *weight* of player i, E is a non-empty set of *resources*, $\Sigma_i \subseteq 2^E \setminus \{\emptyset\}$ is a non-empty *set of strategies* for player i and $\ell_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is the *latency function* of resource $e \in E$. Denote as $\Sigma = \prod_{i\in[n]} \Sigma_i$ the set of all strategy profiles of CG, that is, the set of outcomes which can be realized when each player $i \in [n]$ chooses a strategy in Σ_i. A *strategy profile* $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$ is then a vector of strategies, where, for each $i \in [n]$, $\sigma_i \in \Sigma_i$ denotes the choice of player i in $\boldsymbol{\sigma}$. For a strategy profile $\boldsymbol{\sigma}$ and a resource $e \in E$, the value $n_e(\boldsymbol{\sigma}) = \sum_{i\in[n]:e\in\sigma_i} w_i$ denotes the *congestion* of resource e in $\boldsymbol{\sigma}$, that is, the sum of the weights of all the players choosing e in $\boldsymbol{\sigma}$. The *individual cost* of player i in $\boldsymbol{\sigma}$ is defined as

$c_i(\boldsymbol{\sigma}) = w_i \sum_{e \in \sigma_i} \ell_e(n_e(\boldsymbol{\sigma}))$.

Given a finite space of functions $\mathcal{F} \subseteq \{f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}\}$, let $\mathcal{B}(\mathcal{F}) = \{f_k : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} \mid k \in [r]\}$ be a basis for $\mathcal{F}$ of cardinality r, whose elements (functions) are numbered from 1 to r. We say that CG is defined over $\mathcal{F}$ if, for each $e \in E$, it holds that $\ell_e = \sum_{k \in [r]} v_k^e f_k$, where $v_k^e \in \mathbb{R}$ is a scalar. Throughout the paper, we will impose only ***minimal assumptions*** on $\mathcal{F}$; in particular, we will assume that any $f \in \mathcal{F}$ is non-negative with $f(x) = 0$ if and only if $x = 0$.

For any n-dimensional vector of (positive) weights $\boldsymbol{w} = (w_1, \ldots, w_n)$, we denote with $\mathcal{C}_{\boldsymbol{w}}(\mathcal{F})$ the class of all the weighted congestion games with players' weights induced by $\boldsymbol{w}$ and defined over $\mathcal{F}$. Moreover, for a fixed quadruple $\mathsf{T}_{\boldsymbol{w}} = ([n], \boldsymbol{w}, E, (\Sigma_i)_{i \in [n]})$, called a *congestion model*, the set $\mathcal{C}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F}) = \{\mathsf{CG} \in \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}) \mid \mathsf{CG} = (\mathsf{T}_{\boldsymbol{w}}, (\ell_e)_{e \in E})\}$ is the set of all the weighted congestion games induced by $\mathsf{T}_{\boldsymbol{w}}$ and defined over $\mathcal{F}$. Note that, since for each game $\mathsf{CG} \in \mathcal{C}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F})$ and $e \in E$ there exist r numbers $v_1^e, \ldots, v_r^e$ such that $\ell_e = \sum_{k \in [r]} v_k^e f_k$, it follows that CG can be specified by the pair $(\mathsf{T}_{\boldsymbol{w}}, (v_k^e)_{e \in E, k \in [r]})$. Moreover, it holds that $\mathcal{C}_{\boldsymbol{w}}(\mathcal{F}) = \bigcup_{\mathsf{T}_{\boldsymbol{w}}} \mathcal{C}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F})$. Finally, we denote with $\Sigma(\mathsf{T}_{\boldsymbol{w}})$ the set of strategy profiles induced by the congestion model $\mathsf{T}_{\boldsymbol{w}}$.

A *generalized weighted congestion game* is a pair (CG, α) where $\mathsf{CG} = ([n], (w_i)_{i \in [n]}, E, (\Sigma_i)_{i \in [n]}, (\ell_e)_{e \in E})$ is a weighted congestion game and $\alpha \in \mathbb{R}^{n \times n}$ is an n-dimensional square matrix. Game (CG, α) has the same set of players and strategies of CG, but the *perceived cost* of player i in the strategy profile $\boldsymbol{\sigma}$ is defined as

$$\widehat{c}_i(\boldsymbol{\sigma}) = \sum_{j \in [n]} \alpha_{ij} c_j(\boldsymbol{\sigma}) = \sum_{j \in [n]} \alpha_{ij} w_j \sum_{e \in E : e \in \sigma_j} \ell_e(n_e(\boldsymbol{\sigma})) = \sum_{e \in E} \sum_{k \in [r]} v_k^e f_k(n_e(\boldsymbol{\sigma})) \sum_{j \in [n] : e \in \sigma_j} \alpha_{ij} w_j,$$

where $c_i(\boldsymbol{\sigma})$ is the individual cost that player i experiences in $\boldsymbol{\sigma}$ in the underlying weighted congestion game CG. Note that, when α is the identity matrix, (CG, α) coincides with CG, while, in all the other cases, (CG, α) may not be isomorphic to any weighted congestion game, so that the set of generalized weighted congestion games expands that of weighted congestion games.

Given a strategy profile $\boldsymbol{\sigma}$, a player $i \in [n]$ and a strategy $x \in \Sigma_i$, we denote with $(\boldsymbol{\sigma}_{-i}, x)$ the strategy profile obtained from $\boldsymbol{\sigma}$ when player i changes her strategy from σ_i to x, while the strategies of all the other players are kept fixed. In particular, for any $\epsilon \geq 0$, the perceived cost suffered by player i in $\boldsymbol{\sigma}$ minus $1 + \epsilon$ times the perceived cost suffered by player i in $(\boldsymbol{\sigma}_{-i}, x)$ in a generalized weighted congestion game can be expressed as follows:

$$\begin{aligned}
\widehat{c}_i(\boldsymbol{\sigma}) - (1+\epsilon) \cdot \widehat{c}_i(\boldsymbol{\sigma}_{-i}, x) &= \sum_{j \in [n]} \alpha_{ij} c_j(\boldsymbol{\sigma}) - (1+\epsilon) \sum_{j \in [n]} \alpha_{ij} c_j(\boldsymbol{\sigma}_{-i}, x) \\
&= \alpha_{ii} w_i \left(\sum_{e \in \sigma_i} \ell_e(n_e(\boldsymbol{\sigma})) - (1+\epsilon) \sum_{e \in x} \ell_e(n_e(\boldsymbol{\sigma}_{-i}, x)) \right) \\
&\quad + \sum_{j \in [n] : j \neq i} \alpha_{ij} w_j \left(\sum_{e \in \sigma_i} \ell_e(n_e(\boldsymbol{\sigma})) - (1+\epsilon) \sum_{e \in x} \ell_e(n_e(\boldsymbol{\sigma}_{-i}, x)) \right) \\
&= \alpha_{ii} w_i \left(\sum_{e \in \sigma_i \setminus x} \ell_e(n_e(\boldsymbol{\sigma})) - (1+\epsilon) \sum_{e \in x \setminus \sigma_i} \ell_e(n_e(\boldsymbol{\sigma}) + w_i) \right) \\
&\quad + \sum_{j \in [n] : j \neq i} \alpha_{ij} w_j \left(\sum_{e \in \sigma_i \setminus x} \ell_e(n_e(\boldsymbol{\sigma})) - (1+\epsilon) \sum_{e \in x \setminus \sigma_i} \ell_e(n_e(\boldsymbol{\sigma}) + w_i) \right).
\end{aligned}$$

Hence, we get

$$\widehat{c}_i(\boldsymbol{\sigma}) - (1+\epsilon)\cdot\widehat{c}_i(\boldsymbol{\sigma}_{-i},x) =$$

$$\sum_{e\in\sigma_i\setminus x}\ell_e(n_e(\boldsymbol{\sigma}))\sum_{j\in[n]:e\in\sigma_j}\alpha_{ij}w_j - (1+\epsilon)\sum_{e\in x\setminus\sigma_i}\ell_e(n_e(\boldsymbol{\sigma})+w_i)\left(\alpha_{ii}w_i+\sum_{j\in[n]:e\in\sigma_j}\alpha_{ij}w_j\right). \quad (1)$$

Next two definitions formalize the two concepts of approximate equilibria that we will consider throughout the paper.

Definition 1 *For any $\epsilon\geq 0$, an* **ϵ-approximate coarse correlated equilibrium** *is a probability distribution $\boldsymbol{p}$ defined over Σ such that, for any player $i\in[n]$ and strategy $x\in\Sigma_i$, it holds that*

$$\sum_{\boldsymbol{\sigma}\in\Sigma}p_{\boldsymbol{\sigma}}\cdot\widehat{c}_i(\boldsymbol{\sigma})\leq(1+\epsilon)\sum_{\boldsymbol{\sigma}\in\Sigma}p_{\boldsymbol{\sigma}}\cdot\widehat{c}_i(\boldsymbol{\sigma}_{-i},x),$$

where, for each $\boldsymbol{\sigma}\in\Sigma$, $p_{\boldsymbol{\sigma}}$ is the probability assigned to $\boldsymbol{\sigma}$ by $\boldsymbol{p}$.

Definition 2 *For any $\epsilon\geq 0$, an* **ϵ-approximate pure Nash equilibrium** *is a strategy profile $\boldsymbol{\sigma}$ such that, for any player $i\in[n]$ and strategy $x\in\Sigma_i$, it holds that $\widehat{c}_i(\boldsymbol{\sigma})\leq(1+\epsilon)\cdot\widehat{c}_i(\boldsymbol{\sigma}_{-i},x)$.*

Denote as $\mathsf{PNE}_\epsilon(\mathsf{CG},\alpha)$ and $\mathsf{CCE}_\epsilon(\mathsf{CG},\alpha)$, respectively, the set of ϵ-approximate pure Nash equilibria and ϵ-approximate coarse correlated equilibria of the generalized weighted congestion game (CG,α). It is easy to see that, for any $\epsilon\geq 0$, an ϵ-approximate pure Nash equilibrium $\boldsymbol{\sigma}$ is an ϵ-approximate coarse correlated equilibrium $\boldsymbol{p}$ such that $p_{\boldsymbol{\sigma}}=1$ and $p_{\boldsymbol{\tau}}=0$ for any $\boldsymbol{\tau}\in\Sigma\setminus\{\boldsymbol{\sigma}\}$. So, $\mathsf{PNE}_\epsilon(\mathsf{CG},\alpha)\subseteq\mathsf{CCE}_\epsilon(\mathsf{CG},\alpha)$. Moreover, the sets $\mathsf{PNE}_0(\mathsf{CG},\alpha)$ and $\mathsf{CCE}_0(\mathsf{CG},\alpha)$ coincide with the sets of pure Nash equilibria and coarse correlated equilibria of (CG,α), respectively.

For an n-dimensional *non-null* square matrix $\beta\in\mathbb{R}^{n\times n}_{\geq 0}$ and a player $i\in[n]$, let $\beta\text{-}cost_i:\Sigma\to\mathbb{R}_{>0}$ be the contribution of player i to the definition of the social function which is defined as follows:

$$\beta\text{-}cost_i(\boldsymbol{\sigma})=\sum_{j\in[n]}\beta_{ij}c_j(\boldsymbol{\sigma})=\sum_{e\in E}\sum_{k\in[r]}v_k^e f_k(n_e(\boldsymbol{\sigma}))\sum_{j\in[n]:e\in\sigma_j}\beta_{ij}w_j.$$

Let $\Delta(\Sigma)$ be the set of all the probability distributions defined over Σ. For a $\boldsymbol{p}\in\Delta(\Sigma)$, the *$\beta$-utilitarian social function* is a function $\beta\text{-}\mathsf{SUM}:\Delta(\Sigma)\to\mathbb{R}_{>0}$ such that

$$\begin{aligned}
\beta\text{-}\mathsf{SUM}(\boldsymbol{p}) &= \sum_{i\in[n]}\mathbb{E}_{\boldsymbol{\sigma}\sim\boldsymbol{p}}\left[\beta\text{-}cost_i(\boldsymbol{\sigma})\right]\\
&= \mathbb{E}_{\boldsymbol{\sigma}\sim\boldsymbol{p}}\left[\sum_{i\in[n]}\beta\text{-}cost_i(\boldsymbol{\sigma})\right]\\
&= \sum_{\boldsymbol{\sigma}\in\Sigma}p_{\boldsymbol{\sigma}}\left(\sum_{e\in E}\sum_{k\in[r]}v_k^e f_k(n_e(\boldsymbol{\sigma}))\sum_{i\in[n]}\sum_{j\in[n]:e\in\sigma_j}\beta_{ij}w_j\right)
\end{aligned}$$

and the *β-egalitarian social function* is a function $\beta\text{-}\mathsf{MAX}:\Delta(\Sigma)\to\mathbb{R}_{>0}$ such that

$$\beta\text{-}\mathsf{MAX}(\boldsymbol{p}) = \max_{i\in[n]}\left\{\mathbb{E}_{\boldsymbol{\sigma}\sim\boldsymbol{p}}\left[\beta\text{-}cost_i(\boldsymbol{\sigma})\right]\right\}$$

$$= \max_{i\in[n]} \left\{ \sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \sum_{e\in E} \sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \beta_{ij} w_j \right\}.^{3}$$

Consider the case in which $\boldsymbol{p} \in \Delta(\Sigma)$ is indeed a strategy profile $\boldsymbol{\sigma} \in \Sigma$. When β is the identity matrix, β-SUM (resp. β-MAX) coincides with the sum (resp. the maximum) of the players' individual costs in the underlying weighted congestion game CG, while, when $\beta = \alpha$, β-SUM (resp. β-MAX) coincides with the sum (resp. the maximum) of the players' perceived costs in (CG, α). In general, an infinite variety of social functions can be defined by tuning the choice of matrix β[4]. For a function $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$, we denote with $\boldsymbol{o}$ the *social optimum*, that is, any strategy profile minimizing β-SF. Note that, by the properties of the latency functions and the definition of β[5], it follows that $\beta\text{-}\mathsf{SF}(\boldsymbol{o}) > 0$. The ϵ-approximate coarse correlated price of anarchy of (CG, α) under the social function β-SF is defined as

$$\mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathsf{CG}, \alpha) = \max_{\boldsymbol{p}\in\mathsf{CCE}_\epsilon(\mathsf{CG},\alpha)} \frac{\beta\text{-}\mathsf{SF}(\boldsymbol{p})}{\beta\text{-}\mathsf{SF}(\boldsymbol{o})},$$

while the ϵ-approximate pure price of anarchy of (CG, α) under the social function β-SF is defined as

$$\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathsf{CG}, \alpha) = \max_{\boldsymbol{\sigma}\in\mathsf{PNE}_\epsilon(\mathsf{CG},\alpha)} \frac{\beta\text{-}\mathsf{SF}(\boldsymbol{\sigma})}{\beta\text{-}\mathsf{SF}(\boldsymbol{o})}.$$

For an n-dimensional vector of weights $\boldsymbol{w} = (w_1, \ldots, w_n)$ and a matrix $\alpha \in \mathbb{R}^{n\times n}$, we denote with $\mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha) = \{(\mathsf{CG}, \alpha) : \mathsf{CG} \in \mathcal{C}_{\boldsymbol{w}}(\mathcal{F})\}$ the set of all the generalized weighted congestion games induced by $\boldsymbol{w}$ and α and defined over $\mathcal{F}$. Similarly, for any congestion model $\mathsf{T}_{\boldsymbol{w}}$, one defines the class $\mathcal{C}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$, so as to obtain $\mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha) = \bigcup_{\mathsf{T}_{\boldsymbol{w}}} \mathcal{C}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$. The worst-case ϵ-approximate coarse correlated price of anarchy of the class $\mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha)$ under the social function β-SF is defined as

$$\mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \sup_{(\mathsf{CG},\alpha)\in\mathcal{C}_{\boldsymbol{w}}(\mathcal{F},\alpha)} \mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathsf{CG}, \alpha).$$

Similarly, one defines the worst-case ϵ-approximate pure price of anarchy of the class $\mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha)$ under the social function β-SF.

By $\mathsf{PNE}_\epsilon(\mathsf{CG}, \alpha) \subseteq \mathsf{CCE}_\epsilon(\mathsf{CG}, \alpha)$, it follows that $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) \leq \mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$ for any real value $\epsilon \geq 0$, n-dimensional vector of weights $\boldsymbol{w}$, finite space of function $\mathcal{F}$, pair of matrices $\alpha \in \mathbb{R}^{n\times n}$ and $\beta \in \mathbb{R}^{n\times n}_{\geq 0}$ and function $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$. Throughout the paper, we will also refer to the worst-case ϵ-approximate pure price of anarchy and to the worst-case ϵ-approximate coarse correlated price of anarchy of subsets of $\mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha)$ which are naturally defined by restriction.

We conclude this section with an easy, although crucial result, stating that, independently of which is the adopted social function, both the worst-case ϵ-approximate pure price of anarchy and the worst-case ϵ-approximate coarse correlated price of anarchy of a class of generalized weighted congestion games remain the same even if one restricts to only those games in the given class whose social optimum has social value equal to one[6]. To this aim, for any function $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$

[4]One could even relax the constraint $\beta \in \mathbb{R}^{n\times n}_{\geq 0}$ and allow for negative entries in matrix β as long as $\sum_{i\in[n]} \beta_{ij} \geq 0$ for each $j \in [n]$ and $\sum_{i\in[n]} \beta_{ij} > 0$ for some $j \in [n]$ which still guarantees either $\beta\text{-}\mathsf{SUM}(\boldsymbol{\sigma}) > 0$ and $\beta\text{-}\mathsf{MAX}(\boldsymbol{\sigma}) > 0$ for each $\boldsymbol{\sigma} \in \Sigma$.

[5]From now on, we will always assume that β is a non-null matrix.

[6]Indeed, such a result implicitly holds for the worst-case ϵ-approximate price of anarchy of any kind of equilibrium.

and matrix $\beta \in \mathbb{R}^{n\times n}_{\geq 0}$, let $\overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F},\alpha) \subset \mathcal{C}_{\boldsymbol{w}}(\mathcal{F},\alpha)$ be the subset of all the generalized weighted congestion games induced by $\boldsymbol{w}$ and α and defined over $\mathcal{F}$ such that the social optimum $\boldsymbol{o}$ satisfies β-$\mathsf{SF}(\boldsymbol{o}) = 1$. Similarly, for any congestion model $\mathsf{T}_{\boldsymbol{w}}$, one defines the class $\overline{\mathcal{C}}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F},\alpha)$, so as to obtain $\overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F},\alpha) = \bigcup_{\mathsf{T}_{\boldsymbol{w}}} \overline{\mathcal{C}}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F},\alpha)$.

Lemma 1 *For any real value $\epsilon \geq 0$, n-dimensional vector of weights $\boldsymbol{w}$, finite space of functions $\mathcal{F}$, pair of matrices $\alpha \in \mathbb{R}^{n\times n}$ and $\beta \in \mathbb{R}^{n\times n}_{\geq 0}$ and function $\mathsf{SF} \in \{\mathsf{SUM},\mathsf{MAX}\}$, it holds that $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF},\mathcal{C}_{\boldsymbol{w}}(\mathcal{F},\alpha)) = \mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF},\overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F},\alpha))$ and $\mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF},\mathcal{C}_{\boldsymbol{w}}(\mathcal{F},\alpha)) = \mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF},\overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F},\alpha))$.*

Proof: Fix a congestion model $\mathsf{T}_{\boldsymbol{w}}$, a pair of matrices $\alpha \in \mathbb{R}^{n\times n}$ and $\beta \in \mathbb{R}^{n\times n}_{\geq 0}$ and a function $\mathsf{SF} \in \{\mathsf{SUM},\mathsf{MAX}\}$. The claim directly follows from the fact that, for any game $\mathcal{G} := \left(\mathsf{T}_{\boldsymbol{w}},(v^e_k)_{e\in E,k\in[r]},\alpha\right) \in \mathcal{C}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F},\alpha)$ such that β-$\mathsf{SF}(\boldsymbol{o}) := x > 0$, there always exists a game $\overline{\mathcal{G}} := \left(\mathsf{T}_{\boldsymbol{w}},(\overline{v}^e_k)_{e\in E,k\in[r]},\alpha\right) \in \overline{\mathcal{C}}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F},\alpha)$, obtained by setting $\overline{v}^e_k = v^e_k/x$, such that, for any $\boldsymbol{\sigma} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$, it holds that

$$\sum_{e\in E}\sum_{k\in[r]} v^e_k f_k(n_e(\boldsymbol{\sigma})) \sum_{i\in[n]}\sum_{j\in[n]:e\in\sigma_j} \beta_{ij}w_j = x\sum_{e\in E}\sum_{k\in[r]} \overline{v}^e_k f_k(n_e(\boldsymbol{\sigma})) \sum_{i\in[n]}\sum_{j\in[n]:e\in\sigma_j} \beta_{ij}w_j$$

and that

$$\max_{i\in[n]} \sum_{j\in[n]} \beta_{ij}w_j \sum_{e\in\sigma_j}\sum_{k\in[r]} v^e_k f_k(n_e(\boldsymbol{\sigma})) = x\cdot \max_{i\in[n]} \sum_{j\in[n]} \beta_{ij}w_j \sum_{e\in\sigma_j}\sum_{k\in[r]} \overline{v}^e_k f_k(n_e(\boldsymbol{\sigma})).$$

Moreover, for any $\boldsymbol{\sigma} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$ and $i \in [n]$, it holds that

$$\sum_{e\in E}\sum_{k\in[r]} v^e_k f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij}w_j = x\sum_{e\in E}\sum_{k\in[r]} \overline{v}^e_k f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij}w_j.$$

That is, for any strategy profile $\boldsymbol{\sigma} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$, the social value of $\boldsymbol{\sigma}$ in game $\mathcal{G}$ is equal to x times the social value of $\boldsymbol{\sigma}$ in game $\overline{\mathcal{G}}$, independently of which is the adopted social function. Moreover, for any strategy profile $\boldsymbol{\sigma} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$ and any $i \in [n]$, the perceived cost of player i in $\boldsymbol{\sigma}$ in game $\mathcal{G}$ is equal to x times the perceived cost of player i in $\boldsymbol{\sigma}$ in game $\overline{\mathcal{G}}$. This implies that $\mathcal{G}$ and $\overline{\mathcal{G}}$ have the same set of equilibria (whatever the concept of equilibrium is defined) and that the ratio between any linear combination of the social values of any set of strategy profiles is the same in both games. □

3 The Main Result

Our main result is the proof of the following general theorem.

Theorem 1 *For any real value $\epsilon \geq 0$, n-dimensional vector of weights $\boldsymbol{w}$, finite space of functions $\mathcal{F}$, pair of matrices $\alpha \in \mathbb{R}^{n\times n}$ and $\beta \in \mathbb{R}^{n\times n}_{\geq 0}$ and function $\mathsf{SF} \in \{\mathsf{SUM},\mathsf{MAX}\}$, it holds that $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF},\mathcal{C}_{\boldsymbol{w}}(\mathcal{F},\alpha)) = \mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF},\mathcal{C}_{\boldsymbol{w}}(\mathcal{F},\alpha))$. Moreover, the value $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF},\mathcal{C}_{\boldsymbol{w}}(\mathcal{F},\alpha))$ can always be determined via the primal-dual method.*

Proof: Fix a real value $\epsilon \geq 0$, an n-dimensional vector of weights $\boldsymbol{w}$, a finite space of functions $\mathcal{F}$, a pair of matrices $\alpha \in \mathbb{R}^{n \times n}$ and $\beta \in \mathbb{R}^{n \times n}_{\geq 0}$ and a function $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$. We prove the claim in four steps.

Step 1) Definition of the *representative* congestion model $\mathsf{T}^*_{\boldsymbol{w}}$.

Let $\mathsf{T}^*_{\boldsymbol{w}} = ([n], \boldsymbol{w}, E^*, (\Sigma^*_i)_{i \in [n]})$ be a congestion model such that

1. $\Sigma^*_i = \{\sigma^*_i, o^*_i\}$ for each $i \in [n]$, i.e., each player $i \in [n]$ has exactly two strategies denoted as σ^*_i and o^*_i;
2. the set of resources E^* and the strategies σ^*_i and o^*_i for each $i \in [n]$ are properly defined in such a way that, for each $P, Q \subseteq [n]$, there exists exactly one resource $e(P,Q) \in E^*$ for which it holds that $\{i \in [n] \mid e(P,Q) \in \sigma^*_i\} = P$ and $\{i \in [n] \mid e(P,Q) \in o^*_i\} = Q$. Hence, $|E^*| = 2^n \cdot 2^n = 4^n$.

Intuitively, the representative congestion model $\mathsf{T}^*_{\boldsymbol{w}}$ is defined in such a way that the pair of strategy profiles $\boldsymbol{\sigma}^* = (\sigma^*_1, \ldots, \sigma^*_n)$ and $\boldsymbol{o}^* = (o^*_1, \ldots, o^*_n)$ is able to encompass all possible configurations of congestions that may arise in any pair of strategy profiles and for any congestion model induced by $\boldsymbol{w}$. In particular, the following fundamental property holds.

Property 1 *For any congestion model $\mathsf{T}_{\boldsymbol{w}} = ([n], \boldsymbol{w}, E, (\Sigma_i)_{i \in [n]})$, resource $e \in E$ and pair of profiles $\boldsymbol{\sigma}', \boldsymbol{\sigma}'' \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$, there always exists a resource $\overline{e} \in E^*$ such that $\{i \in [n] \mid e \in \sigma'_i\} = \{i \in [n] \mid \overline{e} \in \sigma^*_i\}$ and $\{i \in [n] \mid e \in \sigma''_i\} = \{i \in [n] \mid \overline{e} \in o^*_i\}$.*

Proof: Fix a congestion model $\mathsf{T}_{\boldsymbol{w}} = ([n], \boldsymbol{w}, E, (\Sigma_i)_{i \in [n]})$, a resource $e \in E$ and pair of profiles $\boldsymbol{\sigma}', \boldsymbol{\sigma}'' \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$. Let $\{i \in [n] \mid e \in \sigma'_i\} := P$ and $\{i \in [n] \mid e \in \sigma''_i\} := Q$. To prove the claim, it suffices choosing $\overline{e} = e(P,Q)$. □

Step 2) Definition of a primal-dual formulation for $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$.

Fix a function $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$. Our aim is to use the optimal solution of a linear program $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ to achieve an upper bound on the worst-case ϵ-approximate pure price of anarchy of any game in $\overline{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$ under the restriction that the latency functions are suitably tuned so as to make $\boldsymbol{\sigma}^*$ the worst ϵ-approximate pure Nash equilibrium and $\boldsymbol{o}^*$ a social optimum (of social value 1). The linear program $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SUM}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ for the β-utilitarian social function is defined as follows.

$$
\begin{aligned}
&\textit{maximize} \sum_{e\in E^*}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{i\in[n]}\sum_{j\in[n]:e\in\sigma_j^*} \beta_{ij} w_j \\
&\textit{subject to} \\
&\sum_{e\in\sigma_i^*\setminus o_i^*}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{j\in[n]:e\in\sigma_j^*} \alpha_{ij} w_j \\
&\quad -(1+\epsilon)\sum_{e\in o_i^*\setminus\sigma_i^*\in E}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma}^*)+w_i)\left(\alpha_{ii}w_i + \sum_{j\in[n]:e\in\sigma_j^*}\alpha_{ij}w_j\right) \le 0, \qquad \forall i\in[n] \\
&\sum_{e\in E^*}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{o}^*)) \sum_{i\in[n]}\sum_{j\in[n]:e\in o_j^*} \beta_{ij} w_j \le 1, \\
&v_k^e \ge 0, \qquad \forall e\in E^*, k\in[r]
\end{aligned}
$$

The first n constraints guarantee that no player can lower her perceived cost of a factor more than $1+\epsilon$ by switching to the strategy she uses in the social optimum $\boldsymbol{o}^*$ (see Equation (1)), while the last constraint normalizes to at most 1 the value β-$\mathsf{SUM}(\boldsymbol{o}^*)$.

The dual program $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SUM}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ is the following (we associate a variable y_i with the ith constraint of the first n ones and a variable γ with the normalizing constraint).

$$
\begin{aligned}
&\textit{minimize } \gamma \\
&\textit{subject to} \\
&\sum_{i\in[n]:e\in\sigma_i^*\setminus o_i^*} y_i f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{j\in[n]:e\in\sigma_j^*}\alpha_{ij}w_j \\
&\quad -(1+\epsilon)\sum_{i\in[n]:e\in o_i^*\setminus\sigma_i^*} y_i f_k(n_e(\boldsymbol{\sigma}^*)+w_i)\left(\alpha_{ii}w_i + \sum_{j\in[n]:e\in\sigma_j^*}\alpha_{ij}w_j\right) \\
&\quad +\gamma f_k(n_e(\boldsymbol{o}^*)) \sum_{i\in[n]}\sum_{j\in[n]:e\in o_j^*}\beta_{ij}w_j \ge f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{i\in[n]}\sum_{j\in[n]:e\in\sigma_j^*}\beta_{ij}w_j, \qquad \forall e\in E^*, k\in[r] \\
&y_i \ge 0, \qquad \forall i\in[n] \\
&\gamma \ge 0
\end{aligned}
$$

Similarly, the linear program $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{MAX}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ for the β-egalitarian social function is defined as follows.

maximize t

subject to

$$\sum_{e\in\sigma_i^*\setminus o_i^*}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{j\in[n]:e\in\sigma_j^*}\alpha_{ij}w_j$$

$$-(1+\epsilon)\sum_{e\in o_i^*\setminus\sigma_i^*\in E}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma}^*)+w_i)\left(\alpha_{ii}w_i+\sum_{j\in[n]:e\in\sigma_j^*}\alpha_{ij}w_j\right)\le 0, \qquad \forall i\in[n]$$

$$\sum_{e\in E^*}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{j\in[n]:e\in\sigma_j^*}\beta_{1j}w_j = t,$$

$$\sum_{e\in E^*}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{j\in[n]:e\in\sigma_j^*}\beta_{ij}w_j \le t, \qquad \forall i\in[n]\setminus\{1\}$$

$$\sum_{e\in E^*}\sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{o}^*)) \sum_{j\in[n]:e\in o_j^*}\beta_{ij}w_j \le 1, \qquad \forall i\in[n]$$

$$v_k^e \ge 0, \qquad \forall e\in E^*, k\in[r]$$

$$t\ge 0$$

Here, again the first n constraints guarantee that no player can lower her perceived cost of a factor more than $1+\epsilon$ by switching to the strategy she uses in the social optimum $\boldsymbol{o}^*$. The next n constraints impose that the maximum value in the social function β-$\mathsf{MAX}(\boldsymbol{\sigma}^*)$ is attained by player 1 (this hypothesis is without loss of generality up to a renumbering of the players) and has value t (which is the objective function to be maximized), while the last n constraints normalizes to at most 1 the value β-$\mathsf{MAX}(\boldsymbol{o}^*)$.

The dual program $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{MAX}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ is the following (we associate variables y_i, z_i and γ_i with the ith constraint of the first, the middle and the last family of n constraints, respectively).

$$
\begin{aligned}
&\textit{minimize} \sum_{i\in[n]} \gamma_i \\
&\textit{subject to} \\
&\sum_{i\in[n]:e\in\sigma^*_i\setminus o^*_i} y_i f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{j\in[n]:e\in\sigma^*_j} \alpha_{ij} w_j \\
&\quad -(1+\epsilon) \sum_{i\in[n]:e\in o^*_i\setminus\sigma^*_i} y_i f_k(n_e(\boldsymbol{\sigma}^*)+w_i) \left(\alpha_{ii} w_i + \sum_{j\in[n]:e\in\sigma^*_j} \alpha_{ij} w_j \right) \\
&\quad +f_k(n_e(\boldsymbol{\sigma}^*)) \sum_{i\in[n]} z_i \sum_{j\in[n]:e\in\sigma^*_j} \beta_{ij} w_j \\
&\quad +f_k(n_e(\boldsymbol{o}^*)) \sum_{i\in[n]} \gamma_i \sum_{j\in[n]:e\in o^*_j} \beta_{ij} w_j \geq 0, && \forall e\in E^*, k\in[r] \\
&\sum_{i\in N} z_i \leq -1 \\
&y_i, z_i, \gamma_i \geq 0, && \forall i\in[n]
\end{aligned}
$$

We stress that, being all the values ϵ, $(w_i)_{i\in[n]}$, $(\alpha_{ij},\beta_{ij})_{i,j\in[n]}$, $n_e(\boldsymbol{\sigma}^*)$ and $n_e(\boldsymbol{o}^*)$ fixed constants in the proposed formulations, $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SUM}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ is a linear program defined over the variables $(v^e_k)_{e\in E^*,k\in[r]}$ and $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{MAX}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ is a linear program defined over the variables $(v^e_k)_{e\in E^*,k\in[r]}$ and t, as needed. Note that, for $\mathsf{SF}\in\{\mathsf{SUM},\mathsf{MAX}\}$, $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ is, in general, under-constrained. In fact, in order to assure that $\boldsymbol{\sigma}^*$ and $\boldsymbol{o}^*$ are the worst ϵ-approximate pure Nash equilibrium and the social optimum, respectively, one should guarantee $\beta\text{-}\mathsf{SF}(\boldsymbol{\sigma}^*) \geq \beta\text{-}\mathsf{SF}(\boldsymbol{\sigma})$ for each other ϵ-approximate pure Nash equilibrium $\boldsymbol{\sigma}\in\Sigma^*$, if any, and $\beta\text{-}\mathsf{SF}(\boldsymbol{o}^*) \leq \beta\text{-}\mathsf{SF}(\boldsymbol{\sigma})$ for each $\boldsymbol{\sigma}\in\Sigma^*$. Moreover, the normalizing constraints have also been relaxed so as to assure $\beta\text{-}\mathsf{SF}(\boldsymbol{o}^*)\leq 1$ rather than $\beta\text{-}\mathsf{SF}(\boldsymbol{o}^*)=1$. Anyway, as we will discuss in the proof of Lemma 2, either removing or relaxing these constraints can only worsen the resulting upper bounds.

The significance of the previously defined pairs of primal-dual formulations is witnessed by the following lemma which states that the value of an optimal solution to $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ provides an upper bound on $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F},\alpha))$.

Lemma 2 *For a fixed* $\mathsf{SF}\in\{\mathsf{SUM},\mathsf{MAX}\}$, *let* $\overline{x}$ *be the value of an optimal solution to* $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ *when this linear problem is not unlimited, otherwise let* $\overline{x}=\infty$. *Then* $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F},\alpha)) \leq \overline{x}$.

Proof: We first show that $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SUM}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ and $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{MAX}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ are both feasible. In fact, fixed an index $k^*\in[r]$, the following solution

$$
v^e_k = \begin{cases} \left(n f_k(w_j) w_j \sum_{i\in[n]} \beta_{ij}\right)^{-1} & \text{if } k=k^* \text{ and } e\in\{e(\{j\},\emptyset), e(\emptyset,\{j\})\} \text{ for some } j\in[n], \\ 0 & \text{otherwise} \end{cases}
$$

is feasible for $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SUM}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ and yields an objective value equal to 1. Similarly, assuming, for instance, that the players are numbered in such a way that $\sum_{i\in[n]} \beta_{1i} \geq \sum_{i\in[n]} \beta_{ji}$ for each $j \in [n] \setminus \{1\}$, the solution with $t = 1$ and

$$v^e_k = \begin{cases} \left(f_k(w_j) w_j \sum_{i\in[n]} \beta_{1i}\right)^{-1} & \text{if } k = k^* \text{ and } e \in \{e(\{j\}, \emptyset), e(\emptyset, \{j\})\} \text{ for some } j \in [n], \\ 0 & \text{otherwise} \end{cases}$$

is feasible for $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{MAX}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ and yields an objective value equal to 1. Hence, for any $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$, exactly one of the two cases included in the claim may occur.

If $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ is unlimited, then $\overline{x} = \infty$ and the claim is trivially true.

So, we can assume that $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ admits an optimal solution of value $\overline{x}$. As we have already observed, $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ may be under-constrained. Nevertheless, recall that we are only interested in an upper bound on the worst-case ϵ-approximate pure price of anarchy of the class $\overline{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$ attainable when the latency functions are suitably tuned so as to make $\boldsymbol{\sigma}^*$ the worst ϵ-approximate pure Nash equilibrium and $\boldsymbol{o}^*$ a social optimum (of social value 1). Let us denote with $\widehat{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$ such a subclass of $\overline{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$. Hence, since once fixed the profiles $\boldsymbol{\sigma}^*$ and $\boldsymbol{o}^*$ any game in $\widehat{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$ can be specified by a particular choice of the values v^e_k, and because the removal or the relaxation of some constraints in a maximization problem can only increase the value of the optimal solution, we obtain that the optimal solution to $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ yields an upper bound on the worst-case ϵ-approximate pure price of anarchy of the class $\widehat{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$. That is, $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \widehat{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)) \leq \overline{x}$. Moreover, since the optimal solution to $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ has value $\overline{x}$, then, by the Strong Duality Theorem, each optimal solution $(\boldsymbol{y}^*, \gamma^*)$ to $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ satisfies $\overline{x} = \gamma^*$. By Property 1, the particular combinatorial structure of the pair $\boldsymbol{\sigma}^*$ and $\boldsymbol{o}^*$ implies that, for any alternative pair of strategy profiles $\boldsymbol{\sigma}$ and $\boldsymbol{o}$ in $\mathsf{T}^*_{\boldsymbol{w}}$, the set of constraints of the dual program $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ is a subset of that of $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$. This implies that any optimal solution $(\boldsymbol{y}, \gamma)$ to $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ must obey $\gamma \leq \gamma^*$. Thus, one can claim that $\gamma^* = \overline{x}$ is indeed an upper bound on the worst-case ϵ-approximate pure price of anarchy of the class $\overline{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$, that is, $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha)) \leq \overline{x}$. Note also that, again by Property 1, for any other congestion model $\mathsf{T}_{\boldsymbol{w}} = ([n], \boldsymbol{w}, E, (\boldsymbol{\Sigma}_i)_{i\in[n]})$, the pair of primal-dual formulations $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ and $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ induced by any pair of strategy profiles $\boldsymbol{\sigma}, \boldsymbol{o} \in \boldsymbol{\Sigma}(\mathsf{T}_{\boldsymbol{w}})$ are such that the set of constraints of $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ is again a subset of that of $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ and this implies that $\gamma^* = \overline{x}$ is even an upper bound on the worst-case ϵ-approximate pure price of anarchy of the whole class $\overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha) = \bigcup_{\mathsf{T}_{\boldsymbol{w}}} \overline{\mathcal{C}}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$, that is, $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) \leq \overline{x}$. □

Step 3) Proof of existence of a game $(\mathsf{CG}, \alpha) \in \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)$ such that $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathsf{CG}, \alpha) = \overline{x}$.

Lemma 3 *For a fixed* $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$*, let* $\overline{x}$ *be the value of an optimal solution to* $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ *when this linear problem is not unlimited, otherwise let* $\overline{x} = \infty$*. Then* $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \overline{x}$.

Proof: Assume that $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ admits a feasible solution $\widehat{\mathsf{SOL}}_{\mathsf{SUM}} = (\widehat{v}^e_k)_{e\in E^*, k\in[r]}$ or $\widehat{\mathsf{SOL}}_{\mathsf{MAX}} = \left((\widehat{v}^e_k)_{e\in E^*, k\in[r]}, \widehat{t}\right)$, both of value $\widehat{x}$, depending on which is the value of SF. Consider the game (CG, α), where $\mathsf{CG} = (\mathsf{T}^*_{\boldsymbol{w}}, (\widehat{v}^e_k)_{e\in E^*, k\in[r]})$ is defined by the representative congestion model $\mathsf{T}^*_{\boldsymbol{w}}$ coupled with the values $(\widehat{v}^e_k)_{e\in E^*, k\in[r]}$. Since, for any $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$, $\widehat{\mathsf{SOL}}_{\mathsf{SF}}$ is feasible for

$\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$, it follows that $\boldsymbol{\sigma}^*$ is an ϵ-approximate pure Nash equilibrium for (CG, α) such that $\beta\text{-}\mathsf{SF}(\boldsymbol{\sigma}^*) = \widehat{x}$. This implies $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) \geq \widehat{x}$ (recall, in fact, that $\beta\text{-}\mathsf{SF}(\boldsymbol{o}^*) \leq 1$).

In the case in which $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ admits an optimal solution $\overline{\mathsf{SOL}}_{\mathsf{SF}}$ of value $\overline{x}$, by the above argument, it follows that $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) \geq \overline{x}$, which, together with Lemma 2, implies the claim. In the case in which $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ is unlimited, then, for any $x \in \mathbb{R}$, there exists a feasible solution $\mathsf{SOL}_{\mathsf{SF}}$ to $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ of value at least x, which implies that, for any $x \in \mathbb{R}$, it holds that $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) \geq x$. □

Step 4) Definition of a primal-dual formulation for $\mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$ and proof of the "Extension Lemma".

Fix a congestion model $\mathsf{T}_{\boldsymbol{w}} = ([n], \boldsymbol{w}, E, (\Sigma_i)_{i\in[n]})$, a probability distribution $\boldsymbol{p} \in \Delta(\Sigma(\mathsf{T}_{\boldsymbol{w}}))$ and a strategy profile $\boldsymbol{o} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$. We define the following primal program $\mathsf{PP}_{\mathsf{CCE}}(\mathsf{SUM}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ for the β-utilitarian social function.

$$
\begin{aligned}
&\textit{maximize} \sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \sum_{e\in E} \sum_{k\in[r]} v^e_k f_k(n_e(\boldsymbol{\sigma})) \sum_{i\in[n]} \sum_{j\in[n]:e\in\sigma_j} \beta_{ij} w_j \\
&\textit{subject to} \\
&\sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \sum_{e\in\sigma_i\setminus o_i} \sum_{k\in[r]} v^e_k f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j \\
&\quad -(1+\epsilon) \sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \sum_{e\in o_i\setminus\sigma_i} \sum_{k\in[r]} v^e_k f_k(n_e(\boldsymbol{\sigma}) + w_i) \left(\alpha_{ii} w_i + \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j \right) \leq 0, \qquad \forall i \in [n] \\
&\sum_{e\in E} \sum_{k\in[r]} v^e_k f_k(n_e(\boldsymbol{o})) \sum_{i\in[n]} \sum_{j\in[n]:e\in o_j} \beta_{ij} w_j \leq 1, \\
&v^e_k \geq 0, \qquad \forall e \in E, k \in [r]
\end{aligned}
$$

The dual program $\mathsf{DP}_{\mathsf{CCE}}(\mathsf{SUM}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ is the following (again, we associate a variable y_i with the ith constraint of the first n ones and a variable γ with the normalizing constraint).

$$
\begin{aligned}
&\textit{minimize } \gamma \\
&\textit{subject to} \\
&\sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \sum_{i\in[n]:e\in\sigma_i\setminus o_i} y_i f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j \\
&\quad -(1+\epsilon) \sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \sum_{i\in[n]:e\in o_i\setminus\sigma_i} y_i f_k(n_e(\boldsymbol{\sigma}) + w_i) \left(\alpha_{ii} w_i + \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j \right) \\
&\quad +\gamma f_k(n_e(\boldsymbol{o})) \sum_{i\in[n]} \sum_{j\in[n]:e\in o_j} \beta_{ij} w_j \geq \sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} f_k(n_e(\boldsymbol{\sigma})) \sum_{i\in[n]} \sum_{j\in[n]:e\in\sigma_j} \beta_{ij} w_j, \qquad \forall e \in E, k \in [r] \\
&y_i \geq 0, \qquad \forall i \in [n] \\
&\gamma \geq 0
\end{aligned}
$$

Similarly, for the β-egalitarian social function, the primal program $\mathsf{PP}_{\mathsf{CCE}}(\mathsf{MAX}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ is defined as follows.

$$
\begin{aligned}
&\textit{maximize } t \\
&\textit{subject to} \\
&\sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \sum_{e\in\sigma_i\setminus o_i} \sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j \\
&\quad -(1+\epsilon) \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \sum_{e\in o_i\setminus\sigma_i} \sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma})+w_i) \left(\alpha_{ii} w_i + \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j\right) \le 0, && \forall i\in[n] \\
&\sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \sum_{e\in E} \sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \beta_{1j} w_j = t, \\
&\sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \sum_{e\in E} \sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \beta_{ij} w_j \le t, && \forall i\in[n]\setminus\{1\} \\
&\sum_{e\in E} \sum_{k\in[r]} v_k^e f_k(n_e(\boldsymbol{o})) \sum_{j\in[n]:e\in o_j} \beta_{ij} w_j \le 1, && \forall i\in[n] \\
&v_k^e \ge 0, && \forall e\in E, k\in[r] \\
&t \ge 0
\end{aligned}
$$

The dual program $\mathsf{DP}_{\mathsf{CCE}}(\mathsf{MAX}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ is the following (again, we associate variables y_i, z_i and γ_i with the ith constraint of the first, the middle and the last family of n constraints, respectively).

$$
\begin{aligned}
&\textit{minimize } \sum_{i\in[n]} \gamma_i \\
&\textit{subject to} \\
&\sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \sum_{i\in[n]:e\in\sigma_i\setminus o_i} y_i f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j \\
&\quad -(1+\epsilon) \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \sum_{i\in[n]:e\in o_i\setminus\sigma_i} y_i f_k(n_e(\boldsymbol{\sigma})+w_i) \left(\alpha_{ii} w_i + \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j\right) \\
&\quad + \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} f_k(n_e(\boldsymbol{\sigma})) \sum_{i\in[n]} z_i \sum_{j\in[n]:e\in\sigma_j} \beta_{ij} w_j \\
&\quad + f_k(n_e(\boldsymbol{o})) \sum_{i\in[n]} \gamma_i \sum_{j\in[n]:e\in o_j} \beta_{ij} w_j \ge 0, && \forall e\in E, k\in[r] \\
&\sum_{i\in N} z_i \le -1 \\
&y_i, z_i, \gamma_i \ge 0, && \forall i\in[n]
\end{aligned}
$$

Again, even though both $\mathsf{PP}_{\mathsf{CCE}}(\mathsf{SUM}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ and $\mathsf{PP}_{\mathsf{CCE}}(\mathsf{MAX}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ may be, in general, under-constrained, by the same arguments used in the discussion of the pairs of primal-dual formulations used for bounding the worst-case ϵ-approximate pure price of anarchy, it follows that,

for each function $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$, the optimal solution to $\mathsf{PP}_{\mathsf{CCE}}(\mathsf{SF}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ yields an upper bound on the worst-case ϵ-approximate coarse correlated price of anarchy of the class $\overline{\mathcal{C}}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$ attainable when $\boldsymbol{p}$ is taken for the worst ϵ-approximate coarse correlated equilibrium and $\boldsymbol{o}$ for the social optimum (of social value 1). Let us denote such a class with $\widehat{\mathcal{C}}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F}, \alpha)$.

The following lemma shows that any upper bound on $\mathsf{PPoA}_{\epsilon}(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\mathsf{T}^*_{\boldsymbol{w}}}(\mathcal{F}, \alpha))$ proved via the primal-dual method automatically extends to $\mathsf{CCPoA}_{\epsilon}(\beta\text{-}\mathsf{SF}, \widehat{\mathcal{C}}_{\mathsf{T}_{\boldsymbol{w}}}(\mathcal{F}, \alpha))$.

Lemma 4 (Extension Lemma) *For any function* $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$*, congestion model* $\mathsf{T}_{\boldsymbol{w}} = ([n], \boldsymbol{w}, E, (\Sigma_i)_{i\in[n]})$*, probability distribution* $\boldsymbol{p} \in \Delta(\Sigma(\mathsf{T}_{\boldsymbol{w}}))$ *and strategy profile* $\boldsymbol{o} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$*, it holds that any feasible solution to* $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ *is also a feasible solution to* $\mathsf{DP}_{\mathsf{CCE}}(\mathsf{SF}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$*.*

Proof: Let $(\boldsymbol{y}^*, \gamma^*)$ be a feasible solution to $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SUM}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$. By Property 1 of the representative congestion model $\mathsf{T}^*_{\boldsymbol{w}}$, it follows that, for any pair of strategy profiles $\boldsymbol{\sigma}, \boldsymbol{o} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$, it holds that

$$\begin{aligned}\sum_{i\in[n]:e\in\sigma_i\setminus o_i} y^*_i f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j \\ - (1+\epsilon) \sum_{i\in[n]:e\in o_i\setminus\sigma_i} y^*_i f_k(n_e(\boldsymbol{\sigma}) + w_i) \left(\alpha_{ii} w_i + \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j\right) \\ + \gamma^* f_k(n_e(\boldsymbol{o})) \sum_{i\in[n]} \sum_{j\in[n]:e\in o_j} \beta_{ij} w_j \geq f_k(n_e(\boldsymbol{\sigma})) \sum_{i\in[n]} \sum_{j\in[n]:e\in\sigma_j} \beta_{ij} w_j \quad (2)\end{aligned}$$

for any $e \in E$ and $k \in [r]$.

Since $p_{\boldsymbol{\sigma}} \geq 0$ for each $\boldsymbol{\sigma} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$, by multiplying inequality (2) for $p_{\boldsymbol{\sigma}}$ and then summing up the obtained inequalities for each $\boldsymbol{\sigma} \in \Sigma(\mathsf{T}_{\boldsymbol{w}})$, we obtain that, for each $e \in E$ and $k \in [r]$, it holds that

$$\begin{aligned}\sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \sum_{i\in[n]:e\in\sigma_i\setminus o_i} y^*_i f_k(n_e(\boldsymbol{\sigma})) \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j \\ - (1+\epsilon) \sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \sum_{i\in[n]:e\in o_i\setminus\sigma_i} y^*_i f_k(n_e(\boldsymbol{\sigma}) + w_i) \left(\alpha_{ii} w_i + \sum_{j\in[n]:e\in\sigma_j} \alpha_{ij} w_j\right) \\ + \gamma^* f_k(n_e(\boldsymbol{o})) \sum_{i\in[n]} \sum_{j\in[n]:e\in o_j} \beta_{ij} w_j \sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} \geq \sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} f_k(n_e(\boldsymbol{\sigma})) \sum_{i\in[n]} \sum_{j\in[n]:e\in\sigma_j} \beta_{ij} w_j. \quad (3)\end{aligned}$$

By $\sum_{\boldsymbol{\sigma}\in\Sigma} p_{\boldsymbol{\sigma}} = 1$, it follows that, for any $e \in E$ and $k \in [r]$, inequality (3) coincides with the relative dual constraint of $\mathsf{DP}_{\mathsf{CCE}}(\mathsf{SUM}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ and this shows that the solution $(\boldsymbol{y}^*, \gamma^*)$ is also feasible for $\mathsf{DP}_{\mathsf{CCE}}(\mathsf{SUM}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$.

A similar argument shows the claim for the case of the social function β-MAX. □

We now have all the ingredients needed to conclude the proof of the theorem.

Fix a function $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$. Assume, first, that $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ is unlimited. Then, by Lemma 3, it holds that $\mathsf{PPoA}_{\epsilon}(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \infty$ which, together with $\mathsf{PPoA}_{\epsilon}(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) \leq \mathsf{CCPoA}_{\epsilon}(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$, immediately implies that

$\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$. By applying Lemma 1, we obtain $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$.

In the case in which $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}^*, \boldsymbol{o}^*)$ admits an optimal solution of value $\overline{x}$, by Lemma 3, it holds that $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \overline{x}$. Moreover, by the Strong Duality Theorem, there exists a feasible solution $(\boldsymbol{y}^*, \gamma^*)$ to $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ of value $\gamma^* = \overline{x}$. Choose an arbitrary game $(\mathsf{CG}, \alpha) \in \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)$ such that $\mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathsf{CG}, \alpha)$ and let $\mathsf{T}_{\boldsymbol{w}}$ be the congestion model defining CG, $\boldsymbol{p}$ be the worst ϵ-approximate coarse correlated equilibrium of (CG, α) and $\boldsymbol{o}$ be the social optimum (of social value 1). By the definition of $\mathsf{T}_{\boldsymbol{w}}$, $\boldsymbol{p}$ and $\boldsymbol{o}$, it follows that the optimal solution to $\mathsf{PP}_{\mathsf{CCE}}(\mathsf{SF}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ has a value of at least $\mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$, which, by the Weak Duality Theorem, implies in turn that any feasible solution to $\mathsf{DP}_{\mathsf{CCE}}(\mathsf{SF}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$ has a value of at least $\mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$. By Lemma 4, it follows that $(\boldsymbol{y}^*, \gamma^*)$ is also a feasible solution to $\mathsf{DP}_{\mathsf{CCE}}(\mathsf{SF}, \mathsf{T}_{\boldsymbol{w}}, \boldsymbol{p}, \boldsymbol{o})$. This implies that $\mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) \leq \gamma^* = \overline{x} = \mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$. Again, by applying Lemma 1, we obtain that $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \mathsf{CCPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$.

It is clear from our discussion that the value $\mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \overline{\mathcal{C}}_{\boldsymbol{w}}(\mathcal{F}, \alpha)) = \mathsf{PPoA}_\epsilon(\beta\text{-}\mathsf{SF}, \mathcal{C}_{\boldsymbol{w}}(\mathcal{F}, \alpha))$ can always be (theoretically) determined via the primal-dual method, that is, by computing the value of the optimal solution of either the primal program $\mathsf{PP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ or the dual one $\mathsf{DP}_{\mathsf{PNE}}(\mathsf{SF}, \mathsf{T}^*_{\boldsymbol{w}}, \boldsymbol{\sigma}, \boldsymbol{o})$ for each function $\mathsf{SF} \in \{\mathsf{SUM}, \mathsf{MAX}\}$, and this concludes the proof (solving the dual program, in particular, requires to determine the minimum value γ^* for which all the $r \cdot 4^n$ possible constraints induced by the $|E^*| = 4^n$ pairs of values yielded by the representative congestion model $\mathsf{T}^*_{\boldsymbol{w}}$ on each of the r components of the latency functions are satisfied). □

4 Conclusions and Open Problems

By introducing the notions of smooth games and robust price of anarchy, Roughgarden [19, 20] showed that the class of congestion games with non-negative and non-decreasing latency functions is tight under the utilitarian social function (see Section 5 of the Appendix for formal definitions). This result has been extended to the class of weighted congestion games by Bhawalkar, Gairing and Roughgarden [3]. By exploiting the primal-dual method we introduced in [4], we have generalized this result along four directions. In fact, our tightness result holds for the class of generalized weighted congestion games, for generalizations of both the utilitarian and the egalitarian social functions, for any non-negative (and possibly decreasing) latency functions and for the approximated version of the price of anarchy.

The fact that two different and seemingly uncorrelated approaches may produce the same type of general results is quite interesting. Understanding whether there is some kind of relationships between them is an intriguing question. Both approaches set up some machinery (smoothness argument vs. primal-dual formulation) allowing for the proof of significant upper bounds on the pure price of anarchy of the games under analysis and then make use of an extension theorem to show that such bounds extend to the coarse correlated price of anarchy as well.

In particular, how is this last step achieved?

Note that the proof of the smoothness argument (see the proof of Lemma 5) requires the definition of smoothness to hold only for any pair of strategy profiles $(\boldsymbol{\sigma}, \boldsymbol{\sigma}')$ such that $\boldsymbol{\sigma}$ is a pure Nash equilibrium and $\boldsymbol{\sigma}'$ is a social optimum. The reason why the definition of smoothness is extended to encompass all possible pairs of strategy profiles is due to the fact that it is indeed the proof of the extension theorem (see the proof of Theorem 2) that asks for such a stronger hypothesis.

Finally, being a coarse correlated equilibrium a particular probability distribution defined over the set of strategy profiles, the notion of smoothness characterizing each profile in the support of any such an equilibrium can be suitably exploited by the linearity of expectation.

In the primal-dual method, instead, the higher degree of generality that is needed to move from pure Nash equilibria up to coarse correlated equilibria is provided by the representative congestion model which imposes that the variables yielding a feasible solution to the dual formulation for the pure price of anarchy have to satisfy any type of "pure dual constraint" that may eventually arise by considering all possible types of configurations of congestions. Then, since it turns out that the dual constraint characterizing the dual formulation for the coarse correlated price of anarchy is indeed a convex combination of a subset of all the possible "pure dual constraints", the extension theorem follows immediately.

Anyway, there is an major difference between the two methods when one aims at showing the tightness of a particular class of games. When adopting the smoothness framework, after having proved that a class of games is (λ, μ)-smooth for a certain pair of parameters λ and μ, one has to show that there exists a game in the class for which the pure price of anarchy is indeed $\frac{\lambda}{1-\mu}$, that is, that the price of anarchy of pure Nash equilibria matches the robust price of anarchy. We stress that this step can be avoided when adopting the primal-dual formulation, since it is directly implied by the Duality Theory (see Lemma 3). In fact, note that the notion of robust price of anarchy, as the best possible upper bound on the pure price of anarchy achievable via the smoothness argument, has no correspondent in the primal-dual method where this bound is implicitly defined by the optimal solution of the pair of primal-dual programs.

By summarizing, our findings seem to reveal that the primal-dual method may be superior to the smoothness framework within the realm of weighted congestion games and their possible generalizations, but, at the same time, the primal-dual method has never been exploited so far outside this realm. Hence, a good starting point would be that of trying to export it to other scenarios of investigation in which the smoothness framework has already been fruitfully applied, such as, for instance, the quantification of the price of anarchy in unrelated scheduling games, valid utility games, opinion formation games and auction theory.

References

[1] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing*, 40(5):1211–1233, 2011.

[2] B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, pp. 57–66, 2005.

[3] K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: price of anarchy, universal worst-case examples, and tightness. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA)*, LNCS 6346, Springer, pp. 17–28, 2010.

[4] V. Bilò. A unifying tool for bounding the quality of non-cooperative solutions in weighted congestion games. In *Proceedings of the 10th Workshop on Approximation and Online Algorithms (WAOA)*, LNCS 7846, Springer, pp. 229–241, 2012.

[5] V. Bilò. On linear congestion games with altruistic social context. In *Proceedings of the 20th International Computing and Combinatorics Conference (COCOON)*, LNCS 8591, Springer, pp. 547–558, 2014..

[6] V. Bilò, A. Fanelli, and L. Moscardelli. On lookahead equilibria in linear congestion games. In *Proceedings of the 9th International Workshop On Internet And Network Economics (WINE)*, LNCS 8289, Springer, pp. 54–67, 2013.

[7] V. Bilò, M. Flammini, and V. Gallotti. On bidimensional congestion games. In *Proceedings of the 19th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, LNCS 7355, Springer, pp. 147–158, 2012.

[8] V. Bilò, M. Flammini, G. Monaco, and L. Moscardelli. Some anomalies of farsighted strategic behavior. *Theory of Computing Systems*, to appear.

[9] V. Bilò and M. Paladini. On the performance of mildly greedy players in cut games. In *Proceedings of the 20th International Computing and Combinatorics Conference (COCOON)*, LNCS 8591, Springer, pp. 513–524, 2014.

[10] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, pp. 67–73, 2005.

[11] G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, LNCS 3669, Springer, pp. 59–70, 2005.

[12] G. Christodoulou, E. Koutsoupias, and P. G. Spirakis. On the performance of approximate equilibria in congestion games. *Algorithmica*, 61(1):116–140, 2011.

[13] B. de Keijzer, Guido Schäfer, A. Anagnostopoulos, and L. Becchetti. Inefficiency of games with social context. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, LNCS 8146, Springer, pp. 219–230, 2013.

[14] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS)*, LNCS 1653, Springer, pp. 404–413, 1999.

[15] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[16] J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Science*, 36(1):48–49, 1950.

[17] M. Rahn and G. Schäfer. Bounding the inefficiency of altruism through social contribution games. In *Proceedings of the 9th International Conference on Web and Internet Economics (WINE)*, LNCS 8289, Springer, pp. 391–404, 2013.

[18] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[19] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, ACM Press, pp. 513–522, 2009.

[20] T. Roughgarden. Intrinsic robustness of the price of anarchy. *Communications of the ACM*, 55(7):116–123, 2012.

[21] L. S. Shapley. The value of n-person games. *Contributions to the theory of games*, Princeton University Press, pp. 31–40, 1953.

Appendix

5 The Smoothness Argument and the Robust Price of Anarchy

Let $\mathcal{G} = \big([n], (\Sigma_i)_{i\in[n]}, (c_i)_{i\in[n]}\big)$ be a cost minimization game defined by the set of players $[n]$, the set of strategies Σ_i and the individual cost function $c_i : \Sigma \to \mathbb{R}_{>0}$ for each player $i \in [n]$. A social function $\mathsf{SF} : \Sigma \to \mathbb{R}_{>0}$ for $\mathcal{G}$ is sum-bounded if, for each $\boldsymbol{\sigma} \in \Sigma$, it holds that $\mathsf{SF}(\boldsymbol{\sigma}) \leq \sum_{i\in[n]} c_i(\boldsymbol{\sigma})$.

Definition 3 (Smoothness) *Given a social function* SF*,* $\mathcal{G}$ *is* (λ, μ)*-smooth under* SF *if, for any two strategy profiles* $\boldsymbol{\sigma}, \boldsymbol{\sigma}' \in \Sigma$*, it holds that* $\sum_{i\in[n]} c_i(\boldsymbol{\sigma}_{-i}, \sigma'_i) \leq \lambda \mathsf{SF}(\boldsymbol{\sigma}') + \mu \mathsf{SF}(\boldsymbol{\sigma})$.

The connection between the notion of smoothness and that of pure price of anarchy is captured by the following lemma.

Lemma 5 (Smoothness Argument) *If* $\mathcal{G}$ *is* (λ, μ)*-smooth under a sum-bounded social function* SF*, with* $\lambda > 0$ *and* $\mu < 1$*, then it holds that* $\mathsf{PPoA}(\mathsf{SF}, \mathcal{G}) \leq \frac{\lambda}{1-\mu}$.

Proof: Let $\boldsymbol{\sigma}$ be any pure Nash equilibrium for $\mathcal{G}$ and $\boldsymbol{o}$ be a social optimum for $\mathcal{G}$ under SF. It holds that

$$\mathsf{SF}(\boldsymbol{\sigma}) \leq \sum_{i\in[n]} c_i(\boldsymbol{\sigma}) \leq \sum_{i\in[n]} c_i(\boldsymbol{\sigma}_{-i}, o_i) \leq \lambda \mathsf{SF}(\boldsymbol{o}) + \mu \mathsf{SF}(\boldsymbol{\sigma})$$

and the claim follows by rearranging the terms. □

The robust price of anarchy is then defined as the best possible upper bound on the pure price of anarchy that can be proved via the smoothness argument. For a game $\mathcal{G}$ and a social function SF, we denote with $\mathcal{A}_{\mathsf{SF}}(\mathcal{G})$ the set of parameters (λ, μ) such that $\mathcal{G}$ is (λ, μ)-smooth under SF.

Definition 4 (Robust Price of Anarchy) *Given a sum-bounded social function* SF*, the robust price of anarchy of* $\mathcal{G}$ *under* SF*, is the value* $\rho_{\mathsf{SF}}(\mathcal{G}) = inf\Big\{\frac{\lambda}{1-\mu} : (\lambda, \mu) \in \mathcal{A}_{\mathsf{SF}}(\mathcal{G})\Big\}$.

The power of the smoothness argument is then stressed by the following extension theorem.

Theorem 2 (Extension Theorem) *For each cost minimization game* $\mathcal{G}$ *and sum-bounded social function* SF *for* $\mathcal{G}$*, it holds that* $\mathsf{CCPoA}(\mathsf{SF}, \mathcal{G}) \leq \rho_{\mathsf{SF}}(\mathcal{G})$.

Proof: Let $\boldsymbol{p}$ be any coarse correlated equilibrium for $\mathcal{G}$ and $\boldsymbol{o}$ be a social optimum for $\mathcal{G}$ under SF. It holds that

$$\sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \cdot \mathsf{SF}(\boldsymbol{\sigma}) \leq \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \sum_{i\in[n]} c_i(\boldsymbol{\sigma}) = \sum_{i\in[n]} \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \cdot c_i(\boldsymbol{\sigma}) \leq \sum_{i\in[n]} \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \cdot c_i(\boldsymbol{\sigma}_{-i}, o_i)$$

$$= \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \sum_{i\in[n]} c_i(\boldsymbol{\sigma}_{-i}, o_i) \leq \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \left(\lambda \mathsf{SF}(\boldsymbol{o}) + \mu \mathsf{SF}(\boldsymbol{\sigma})\right) = \mu \sum_{\boldsymbol{\sigma}\in\boldsymbol{\Sigma}} p_{\boldsymbol{\sigma}} \cdot \mathsf{SF}(\boldsymbol{\sigma}) + \lambda \mathsf{SF}(\boldsymbol{o})$$

and the claim follows by rearranging the terms. □

Let $\mathcal{C}$ be a class of cost minimization games and $\widehat{\mathcal{C}} \subseteq \mathcal{C}$ be the subclass of the games in $\mathcal{C}$ which admit at least one pure Nash equilibrium. Given a social function SF, we denote with $\mathcal{A}_{\mathsf{SF}}(\mathcal{C})$ the set of parameters (λ, μ) such that each game $\mathcal{G} \in \widehat{\mathcal{C}}$ is (λ, μ)-smooth under SF.

Definition 5 (Tight Class of Games) *A class of games $\mathcal{C}$ is tight under the social function SF if it holds that $sup_{\mathcal{G}\in\widehat{\mathcal{C}}} \mathsf{PPoA}(\mathsf{SF}, \mathcal{G}) = inf_{(\lambda,\mu)\in\mathcal{A}(\mathcal{C})} \frac{\lambda}{1-\mu}$.*